

Addendum to “A Market Framework for Eliciting Private Data”

Bo Waggoner¹, Rafael Frongillo², and Jacob Abernethy³

¹Harvard SEAS. bwaggoner@fas.harvard.edu

²University of Colorado. raf@colorado.edu

³University of Michigan. jabernet@umich.edu

Abstract

In Waggoner et al. [2015], we proposed an elaboration on prediction markets that preserves differential privacy of transactions. This addendum gives a simpler, prediction-market focused construction and proofs. It also shows how to recover forms of a bounded worst-case loss guarantee by introducing a transaction fee.

1 Overview of Prediction Markets

To begin, we briefly review cost-function based prediction markets and introduce the notation we use for the rest of the addendum. Recall that the goal of a prediction market is to aggregate information about a future event Z , a random variable which has a finite set of possible outcomes \mathcal{Z} . The market maker specifies a set of *securities* \mathcal{F} , which are functions $\phi : \mathcal{Z} \rightarrow \mathbb{R}$. Traders sequentially arrive and interact with the market maker by offering to buy *shares* of one or more securities. After trading concludes, the outcome Z is observed, and each share in ϕ pays off $\phi(Z)$.

In the classic market formulation we focus on here, the securities are “Arrow-Debreu” securities: There is one security ϕ_z per $z \in \mathcal{Z}$, with $\phi_z(Z) = 1$ if $z = Z$ and $\phi_z(Z) = 0$ otherwise. In other words, a trader will receive a payout of 1 for each share she owns in the observed outcome. Naturally, a trader who believes the probability of z is $p(z)$ expects to make money on average for buying a share in ϕ_z if the price is anything less than $p(z)$. Hence, the price of each security is interpreted as a prediction of the probability of the associated outcome. (Note that traders may “short sell” by purchasing a negative number of shares in ϕ , receiving a payoff of $-\phi(Z)$ for each share.) The market is called “complete” because there are $|\mathcal{Z}|$ securities, one for each outcome, whose prices determine a probability distribution on \mathcal{Z} . Note that it also suffices to have only $|\mathcal{Z}| - 1$ securities, with current prices implicitly giving a prediction for the chance of the remaining outcome; a trader who thinks that outcome is more likely can express this by short-selling all of the available securities.

The main design question is how to determine the cost of each purchase, in a manner which may depend on the prior history of the market. In a *cost-function based* prediction market [Abernethy et al., 2013], there is a *market state* $\theta \in \mathbb{R}^{\mathcal{F}}$, which we can think of as a vector giving the total number of shares purchased in each security so far. The market maker initially chooses a convex “cost function” $C : \mathbb{R}^{\mathcal{F}} \rightarrow \mathbb{R}$ which acts as a potential function. When a trader arrives and purchases a share vector (also called a *bundle*) $d\theta$, the market-maker charges $C(\theta + d\theta) - C(\theta)$. The trader will receive $d\theta(Z)$ when Z is observed.

2 Private Markets

A participant in the market who makes a trade $d\theta$ may not want this trade revealed to others, particularly if it is based on private data or information. However, the markets' updates to prices might reveal information about $d\theta$. Here, we modify the classic prediction market to guarantee privacy as formalized by ϵ -*differential privacy*, to be defined shortly. The idea is to add noise to participants' trades so that their true trade cannot be inferred. The guarantee obtained is that no observers or other participants, even if they all collude to share information, can learn “very much” about one participant's trade from the outputs of the market.¹

We will focus on classic, complete cost function markets with d securities.

Two key properties of prediction markets are incentive alignment and bounded worst-case loss of the market maker. We will first show how to satisfy incentive alignment while achieving a nontrivial accuracy guarantee under differential privacy. Cummings et al. [2016] showed that a constant worst-case loss bound (regardless of the number of traders) cannot be achieved while satisfying differential privacy. However, we will show that by introducing a transaction fee, a form of bounded worst-case loss can be restored while still maintaining approximate incentive alignment and accuracy.

Outline. In Section 2.1, we overview a key relevant concept: the *price sensitivity* (a variant of liquidity) of the market. The discussion, it is hoped, will be helpful in giving perspective on the following formal results.

In Section 2.2, we show how to construct a privacy-preserving market with good accuracy and incentive properties. The key challenge is that enough noise must be added to each trade so as to completely obscure it, yet the total noise must be much less than the total trades to allow information to aggregate. We use a technique from the differential privacy literature to overcome this challenge. The key parameter turns out to be the *price sensitivity* of the cost function: how rapidly prices adapt to trades.

In Section 2.3, we show how introducing a transaction fee allows us to recover bounded worst-case loss. The main idea is to relate the price sensitivity to the number of arriving traders.

2.1 Key concept: price sensitivity

The price sensitivity of a cost function C is a measure of how quickly prices respond to trades. This is essentially the same notion as “liquidity” discussed in Abernethy et al. [2013, 2014]. Formally, the *price sensitivity* λ of C is the supremum of the operator norm of the Hessian of C , with respect to the ℓ_1 norm.² The key implication is a Lipschitz condition that if L shares are purchased, then the prices change by no more than λL . A very small price sensitivity means that prices react very slowly to trades.

Price sensitivity is also directly related to the worst-case loss guarantee of the market, as follows. Those familiar with market scoring rules may recall that with scoring rule S , the loss can be bounded by (a constant times) the largest possible score. Hence, scaling S by a factor $\frac{1}{\lambda}$ immediately scales the loss bound by $\frac{1}{\lambda}$ as well. Recall that S is defined by a convex function G , the convex conjugate of C . Scaling S by $\frac{1}{\lambda}$ is equivalent to scaling G by $\frac{1}{\lambda}$. By standard results in convex analysis,

¹This kind of guarantee is sometimes called *joint differential privacy*. We will formally define the precise privacy guarantee below.

²For convenience we will assume C is twice differentiable, though this is not necessary.

this is equivalent to transforming C into $C_\lambda(\theta) = \frac{1}{\lambda}C(\lambda\theta)$, an operation known as the perspective transform. This in turn scales the price sensitivity by λ by the properties of the Hessian.

Convention: normalized, scaled C . In the remainder of the paper, we will suppose that we start with some C' whose price sensitivity equal to 1 and worst-case loss bounded by B for some constant B . We will always use the cost function $C(\cdot) = \frac{1}{\lambda}C'(\lambda\cdot)$, where λ will be chosen in context. As discussed above, C has price sensitivity at most λ and a worst-case loss bound of B/λ .

This assumption is without loss of generality, as there are known cost functions for which price sensitivity and worst-case loss are bounded by constants.

Relation to privacy. To achieve privacy we will add random noise to the market state and publish the associated prices. For a fixed trade size (we will assume all trades are bounded), differential privacy specifies a fixed level of noise that, when added to the market state θ , is sufficient to protect privacy. Given this noise level, the price sensitivity λ mediates between accuracy of prices and worst-case loss of the market maker. A very small λ means that prices react very slowly to changes in the market state. This means that the fixed amount of noise has a small effect on the overall prices, leading to a high level of accuracy. However, the worst-case loss scales by $1/\lambda$.

The goal will be to set a desired, constant level of accuracy—for instance, all prices are within α of what they would be without the noise—and ask how large one can make λ , equivalently, how small the market maker’s loss can be. We will show that, using a technique from differential privacy, it can suffice to have λ only shrinking as $\frac{1}{(\log T)^2}$ when there are T traders, leading to a worst-case loss bound of $O((\log T)^2)$.

This tradeoff is closely related to a measure of expressiveness of the market, namely, how many knowledgeable traders must arrive before the prices have moved to reflect their beliefs. In fact, by definition of price sensitivity, this number must be on the order of $\frac{1}{\lambda}$, as each trader can only move the prices by λ . So our result can also be interpreted as saying that, when the market is designed to protect privacy of T arriving traders, it nonetheless only requires $(\log T)^2$ arrivals before prices accurately reflect trader beliefs.

Worst-case loss and transaction fee. The private market causes two sources of unbounded loss for the market maker. The first is from traders betting against the random noise introduced to protect privacy. As far as we know, this could cause a loss of $\Omega(T)$ with T arrivals. However, a transaction fee can be chosen to exactly balance the expected profit from this type of arbitrage. We will show that this fee is still small enough to allow for very accurate prices.³ This transaction fee restores the worst-case loss guarantee to the inverse of the price sensitivity just as in a non-private market.

Unfortunately, due to the privacy-preserving market construction, we must set $\lambda = \Theta(1/(\log T)^2)$ so that $1/\lambda = \Theta((\log T)^2)$, which is a much better loss guarantee than $\Omega(T)$, but nonetheless increasing with T . Moreover, a transaction fee alone cannot mitigate the worst-case loss of $1/\lambda$, by the following intuition. Traders will only participate if they expect to make more money from the trade than the transaction fee, so the transaction fee must be smaller than the desired accuracy level of prices. But after about $1/\lambda$ arrivals, prices have moved to reflect trader beliefs, so in the

³For instance, if the current price of a security is 0.49 and a trader believes the true price should be 0.50, she will only purchase a share if the transaction fee $c < 0.01$. (Recall that, for privacy reasons, we are also limiting each trade to a fixed size, say, one share.)

worst case are an approximately optimal prediction. Each arrival expects to make a constant-sized profit, *even after subtracting the transaction fee*, which means total market loss must still be on the order of $1/\lambda$ even counting the transaction fee gains.

Obtaining bounded worst-case loss via adaptive liquidity. Eventually, we will show that one can construct a version of prediction markets, intuitively sharing their properties, with worst-case loss bounded by a constant independent of T . The key will be the following observation. After the prices have moved to the optimal prediction, if more traders arrive, then they can only improve market maker profit: the worst-case loss does not change, but transaction fees are collected. In fact, in our case, if all T participants arrive, then the total profit from transaction fees is $\Theta(T)$ while the worst-case loss from the market is $O((\log T)^2)$.

We can leverage this to achieve a bounded worst-case loss with an “adaptive-liquidity” approach. Begin by setting $T^{(1)} = O(1)$, with $\lambda^{(1)}$ on the order of $1/(\log T^{(1)})^2 = \Theta(1)$, and run a private market for $T^{(1)}$ participants. If fewer than $T^{(1)}$ participants show up, the worst-case loss is order $1/\lambda^{(1)}$, a constant.

If all $T^{(1)}$ participants arrive, then (for the right choice of constants) the market has actually turned a profit $\Omega(T^{(1)})$ from the transaction fees. Now set up a private market for $T^{(2)} \gg T^{(1)}$ traders with $\lambda^{(2)}$ on the order of $1/(\log T^{(2)})^2$. If fewer than $T^{(2)}$ participants arrive, the worst-case loss is order $1/\lambda^{(2)}$. However, by design, we will have this loss smaller than the $\Omega(T^{(1)})$ profit from the previous market. So total worst-case loss remains bounded by a constant.

If all $T^{(2)}$ participants arrive, then again this market has turned a profit, which can be used to completely offset the worst-case loss of the next market, and so on.

By adaptively changing the price sensitivity at each round, we will be able to maintain an accuracy guarantee where, as the market gets larger and larger, the impact of the privacy noise is smaller and smaller. An interesting direction for future work would be to replace the graduated approach here with the continuous liquidity adaptation of Abernethy et al. [2014].

2.2 Constructing private markets

General approach. In a private market, the designer chooses an initial “true” market state θ^0 (for convenience, we will assume $\theta^0 = 0$) and announces a “published” market state $\hat{\theta}^0 = 0$. When participant $t = 1, \dots, T$ arrives and requests trade $d\theta^t$, the market maker updates the true market state $\theta^t = \theta^{t-1} + d\theta^t$, but does not reveal the true state to anyone. Instead, the market maker announces the published market state $\hat{\theta}^t$, which is some randomized function of all trades and published market states so far. We assume that $\|d\theta^t\|_1 \leq 1$, that is, each participant can buy or sell at most one “total” share.⁴

Differential privacy. The market mechanism can be viewed as a randomized function M that takes as input a list of trades $\vec{d\theta} = d\theta^1, \dots, d\theta^T$ and outputs a list of published market states $\hat{\theta}^0, \dots, \hat{\theta}^T$. We will call it (ϵ, δ) -*differentially private* if changing a single participant’s trade does not change the distribution on outputs much: if for all $\vec{d\theta}$ and $\vec{d\theta}'$ differing only in one entry, and for all (measurable) sets of possible outputs S ,

$$\Pr \left[M \left(\vec{d\theta} \right) \in S \right] \leq e^\epsilon \Pr \left[M \left(\vec{d\theta}' \right) \in S \right] + \delta.$$

⁴One could also modify our approach to allow arbitrarily large trades, but this would also require adding proportionally large noise in order to continue to preserve privacy.

It is reasonable to treat ϵ as a constant whose size controls the privacy guarantee, such as $\epsilon = 0.01$. Meanwhile, δ is normally preferred to be vanishingly small or 0, as a mechanism can leak the private information of all individuals with δ probability and still be (ϵ, δ) -differentially private.

To be careful, we note that the market’s “full output” also includes that it sends each participant their payoff. However, this payoff is a function only of the public noisy market states and of that participant’s trade. The payoff is assumed to be sent privately and separately, unobservable by any other party. By the post-processing property of differential privacy, a trader’s (ϵ, δ) -privacy guarantee continues to hold regardless of how the published market states are combined with any side information, even including the full list of all other participant’s trades.

Tool 1: Laplace noise. Imagine that the market could first collect all T trades simultaneously, then sum them and publish some $\hat{\theta}^T$, a noisy version of the market state $\theta^T = \sum_{t=1}^T d\theta^t$.

In this scenario, there is only one output $\hat{\theta}^T$ instead of a whole list of outputs $\hat{\theta}^1, \dots$. The standard, simplest solution to protecting privacy would be to take the true sum θ^T and add noise from the Laplace distribution, where a $Lap(b)$ variable has probability density $x \mapsto \frac{1}{2b}e^{-|x|/b}$. In particular, we let z be a vector where each entry is an independent $Lap(2/\epsilon)$, with length equal to the number of securities, and let $\hat{\theta}^T = \theta^T + z$.⁵

Then releasing $\hat{\theta}^T$ is $(\epsilon, 0)$ -differentially private. Note that it also satisfies a good accuracy guarantee, as the amount of noise required does not scale with T ; so with enough participants, this mechanism becomes a very accurate indication of the “average” trade while still preserving privacy.

Tool 2: continual observation technique. Unfortunately, the above solution is not sufficient because our market must publish a market state at each time step. One naive approach is to apply the above solution independently at each time step, *i.e.* produce each $\hat{\theta}^t = \theta^t + z^t$ where z^t contains independent Laplace noise. The problem is that each step reveals more information about a trade, for instance, $d\theta^1$ participates in T separate publications. To continue preserving privacy, each z^t must have a much larger variance (a smaller parameter), which makes the published market states very inaccurate.

A second naive approach is to add noise to each $d\theta^t$ just once, producing $\hat{d\theta}^t = d\theta^t + z^t$. Then set $\hat{\theta}^t = \sum_{s=1}^t \hat{d\theta}^s$. The benefit to this approach is that it can re-use the noisy z^t variables across time steps, rather than re-drawing new noise each time. The problem is that, while each z^t is small in magnitude, there are many of them; for example, the final $\hat{\theta}^T$ contains T pieces of noise, which add up to a very inaccurate estimate of the true market state. This contrasts with the first naive approach, in which each publication only includes one piece of noise, but that piece of noise is very large.

The idea of the “continual observation” technique, pioneered by Dwork et al. [2010] and Chan et al. [2011], is to strike a balance between these extremes by re-using noise a limited number of times while also keeping each piece of noise small. Roughly, each publication $\hat{\theta}^t$ will include a logarithmic (in t) number of pieces of noise, each of which is only “logarithmically large”.

Definition 1. The *private market mechanism* uses a cost function C with parameter λ . At each time t participant t arrives and proposes trade $d\theta^t$ satisfying $\|d\theta^t\|_1 \leq 1$. At most T participants

⁵This is not the most sophisticated solution to preserving privacy of a vector, but it is simple and sufficient for our purposes.




Figure 1: Picturing the continual observation technique for preserving privacy [Dwork et al., 2010, Chan et al., 2011]. Each $d\theta^t$ is a trade. The true market state at t is $\theta^t = \sum_{j=1}^t d\theta^j$ and the goal is to release a noisy version $\hat{\theta}^t$. The goal is to release, at each time step t , a noisy version $\hat{\theta}^t$ of θ^t . Each arrow originates at t , points backwards to $s(t)$, and is labeled with independent Laplace noise vector z^t . Now $\hat{\theta}^t = \theta^t + z^t + z^{s(t)} + z^{s(s(t))} + \dots$. In other words, the noise added at t is a sum of noises obtained by following the arrows all the way back to 0. There are two key properties: Each t has only $\log T$ arrows passing above it, and each path backwards takes only $\log T$ jumps.

may arrive. Let $\theta^t = \sum_{j=1}^t d\theta^j$. At each time t , the mechanism publishes market state

$$\hat{\theta}^t := \theta^t + z^t + z^{s(t)} + z^{s(s(t))} + \dots + z^0,$$

where $z^0 = 0$ and z^t has independent noise on each coordinate drawn $Lap(2\lceil \log T \rceil / \epsilon)$. Here $s(t)$ is defined by writing the integer t in binary, then flipping the rightmost “one” bit to zero. Hence $s(0) = 0$ and $s(t) < t$ for all $t > 0$. Participant t is charged $C(\hat{\theta}^t + d\theta^t) - C(\hat{\theta}^t)$. When outcome Z occurs, she is paid $d\theta^t(Z)$.

A convenient notation is to let

$$\hat{\theta}^{s(t):t} := \left(\sum_{j=s(t)+1}^t d\theta^j \right) + z^t.$$

Then we can define the mechanism recursively as

$$\begin{aligned} \hat{\theta}^t &= \theta^t + z^t + z^{s(t)} + z^{s(s(t))} + \dots + z^0 \\ &= \hat{\theta}^{s(t):t} + \hat{\theta}^{s(t)}. \end{aligned}$$

Remark. Notice that λ has no impact on the construction of the market, amount of noise to add, etc. Intuitively, this is because the market is defined entirely in “share space”, while price sensitivity relates shares to prices. We will not need to discuss λ until we discuss accuracy of the prices, which is irrelevant to the proof of privacy.

Theorem 1 (Privacy). *Assuming that all trades satisfy $\|d\theta^t\|_1 \leq 1$, the private mechanism is ϵ -differentially private in the trades $d\theta^1, \dots, d\theta^T$ with respect to the output $\hat{\theta}^1, \dots, \hat{\theta}^T$.*

Proof. Suppose that the market published all partial sums it uses, i.e. $\hat{\theta}^{s(t)+1:t}$ for all time steps t . Because these values completely determine the market outputs $\hat{\theta}^t$ s, it suffices to show that this would be ϵ -differentially private.

The idea is to treat each publication of the form $\hat{\theta}^{s(t)+1:t}$ as a separate mechanism, which has a guarantee of $(\epsilon/\lceil \log T \rceil)$ -differential privacy. We then show that any one trade $d\theta^t$ participates in at most $\lceil \log T \rceil$ of these mechanisms. These two claims imply the result because, by the composition property of differential privacy [Dwork and Roth, 2014], each trade is therefore guaranteed $\lceil \log T \rceil \cdot (\epsilon/\lceil \log T \rceil) = \epsilon$ -differential privacy.

First, we claim that each publication $\hat{\theta}^{s(t)+1:t}$ preserves $\epsilon/\lceil \log T \rceil$ differential privacy of each trade $d\theta^{t'}$ that participates, i.e. with $s(t) < t' \leq t$. To show this, consider two possible trades $d\theta_1^{t'}, d\theta_2^{t'}$ each of norm at most 1 and let $\theta_1^{s(t)+1:t}, \theta_2^{s(t)+1:t}$ be the resulting true partial sums. If we draw each coordinate's noise independently from a Laplace distribution with parameter c , then for any vector y , the ratio of the density functions is (recall d is the dimension of each θ^t)

$$\begin{aligned} \frac{\Pr[\hat{\theta}_1^{s(t)+1:t} = y]}{\Pr[\hat{\theta}_2^{s(t)+1:t} = y]} &= \prod_{i=1}^d \frac{e^{-|\theta_1^{s(t)+1:t}(i) - y(i)|/c}}{e^{-|\theta_2^{s(t)+1:t}(i) - y(i)|/c}} \\ &\leq \prod_{i=1}^d e^{|\theta_1^{s(t)+1:t}(i) - \theta_2^{s(t)+1:t}(i)|/c} \\ &= e^{\|d\theta_1^{t'} - d\theta_2^{t'}\|_1/c} \\ &\leq e^{2/c}. \end{aligned}$$

(In the first line, we have canceled out the normalizing constants of the Laplace density function; in the last, we used the bounded-norm assumption on trades.) Now, we have the noise parameter $c = 2\lceil \log T \rceil/\epsilon$ by design, showing that the ratio is at most $e^{\epsilon/\lceil \log T \rceil}$, completing the claim.

Second, we claim that each trade $d\theta^{t'}$ participates in at most $\lceil \log T \rceil$ different partial sums $\hat{\theta}^{s(t)+1:t}$. To show this, we only need to count the time steps t where $s(t) < t' \leq t$, in other words, integers $t \geq t'$ where zeroing the rightmost “one” bit gives a number less than t' .

Without loss of generality, the binary expansion of t is $b_m b_{m-1} \dots b_j 10 \dots 0$ for some m, j and then $s(t)$ has expansion $b_m b_{m-1} \dots b_j 00 \dots 0$. Hence the condition $s(t) < t' \leq t$ implies that the binary expansion of t matches that of t' from bits m to j , then has a one at bit $j-1$, and has zeroes at all lower-order bits. Since m is fixed for t' , this can only happen once for each j , or at most m total times; and $m \leq \lceil \log T \rceil$ because $t' \leq T$. \square

Lemma 1 (Accuracy of share vector). *In the private mechanism with d securities and T time steps, let $\beta = \max_{1 \leq t \leq T} \|\theta^t - \hat{\theta}^t\|_1$. Then with probability $1 - \gamma$,*

$$\beta \leq \frac{4\sqrt{2}d \log \lceil T \rceil}{\epsilon} \ln \left(\frac{2Td}{\gamma} \right).$$

Proof. For each t , each coordinate i of $\theta^t - \hat{\theta}^t$ is the sum of at most $\lceil \log T \rceil$ independent variables distributed $\text{Lap}(2\lceil \log T \rceil/\epsilon)$. We will choose β such that each coordinate's absolute value exceeds

β with probability $\frac{\gamma}{Td}$; there are d coordinates per time step and T time steps, so a union bound gives the result.

Choose β such that, if Y is the sum of $k = \lceil \log T \rceil$ independent $Lap(b)$ variables with $b = 2\lceil \log T \rceil/\epsilon$ variables, then

$$\Pr[|Y| > \beta] = \gamma'.$$

A concentration bound for the sum of k independent $Lap(b)$ variables, Corollary 12.3 of Dwork and Roth [2014]⁶ gives

$$\beta \leq 2\sqrt{2}b \ln \frac{2}{\gamma'}.$$

Now choose $\gamma' = \frac{\gamma}{Td}$. To recap, each $|\theta^t(i) - \hat{\theta}^t(i)| \leq \beta$ except with probability $\gamma' = \frac{\gamma}{Td}$, hence by a union bound this holds for all t, i except with probability γ , hence $\|\theta^t - \hat{\theta}^t\|_1 \leq d\beta$ except with probability γ . \square

As mentioned above, the previous results (Theorem 1 and Lemma 1) do not depend on λ at all, because they do not mention the prices. We now ask what a “reasonable” choice of λ can be so that the prices are interpretable as predictions, i.e. the prices are “accurate”.

Theorem 2 (Accuracy of prices). *In the private mechanism, let $p^t = \nabla C(\theta^t)$ and let $\hat{p}^t = \nabla C(\hat{\theta}^t)$. Then to satisfy $\|p^t - \hat{p}^t\|_1 \leq \alpha$ for all t , except with probability γ , it suffices for the price sensitivity to be*

$$\lambda^* \leq \frac{\alpha \epsilon}{4\sqrt{2}d \lceil \log T \rceil \ln(2Td/\gamma)}.$$

Proof. By definition of λ , we have

$$\begin{aligned} \|p^t - \hat{p}^t\|_1 &\leq \lambda \|\theta^t - \hat{\theta}^t\|_1 \\ &\leq \lambda \frac{4\sqrt{2}d \log \lceil T \rceil}{\epsilon} \ln \left(\frac{2Td}{\gamma} \right) \end{aligned} \tag{1}$$

for all t except with probability γ , by Lemma 1. We now just choose λ so that (1) $\leq \alpha$. \square

2.3 Transaction fee and bounded worst-case loss

Unfortunately, for the private prediction market described above, we do not have a better bound than $O(T)$ on the market maker’s loss. Consider an event that is guaranteed to happen and participants who always buy a share of a security in that event; after each trade, the noise added has a half chance to reduce the prices, allowing the next trader to purchase a share at a discount relative to the “true” market prices. These kinds of arbitrages then accumulate. Indeed, Cummings et al. [2016] shows that no prediction market of the form defined above can simultaneously be private and bound its loss by a constant.

We will take two stabs at overcoming this impossibility as much as possible by introducing a small transaction fee c for each trade.

1. We will show that the loss of the private market becomes $O(1/\lambda^*) = O(\log(T)^2)$. This is viewed as a positive result because the worst-case loss is growing quite slowly in the total number of participants, and moreover matches the fundamental “informational” worst-case loss one expects when price sensitivity is λ^* .

⁶In the parameters of that Corollary, we choose $\nu = b\sqrt{\ln(2/\gamma')}$ as we will have $\ln(2/\gamma') > k$.

2. We will design a market consisting of a series of the above private markets, each with adaptively decreasing price sensitivity. This construction will have worst-case loss bounded by an absolute constant independent of the number of arriving traders T (but depending on the other parameters).

The idea is to make the fee large enough to cover expected losses from “noise arbitrage”, where participants only correct for the privacy-preserving noise of the mechanism; but keep it small enough to encourage “information trading”, where participants with beliefs significantly different than the current prices expect to still make money by trading. Thus, the accuracy and incentive alignment properties of the private market are essentially preserved despite the fee.

Theorem 3. *With transaction fee $c = \alpha$, the private market with $\lambda = \lambda^*$ (Theorem 2) has worst-case loss bounded by $\frac{B}{\lambda} = O((\log T)^2)$, fixing $d, \alpha, \gamma, \epsilon$.*

Furthermore, it satisfies the same accuracy guarantee of Theorem 2, namely $\|p^t - \hat{p}^t\|_1 \leq \alpha$ for all t except with probability γ .

Furthermore, its incentive guarantee is that, if participant t with belief p has $\|p - \hat{p}^t\|_\infty \geq 2\alpha$, then t has the incentive to purchase a bundle (again except with prob. γ).

Proof. For the incentive guarantee: a participant’s expected profit from purchasing $d\theta$ is

$$\text{profit} = \langle d\theta, p \rangle + C(\hat{\theta}^t) - C(\hat{\theta}^t + d\theta)$$

Intuitively, the price paid (difference in cost function) is very closely approximated by $\langle \nabla C(\hat{\theta}^t), d\theta \rangle$. Because $\nabla C(\hat{\theta}^t) = \hat{p}^t$, this approximation would immediately prove it. Formally, we will use the small price sensitivity (so $\nabla C(\hat{\theta}^t + d\theta)$ is also close to \hat{p}^t).

$$\begin{aligned} \text{profit} &= \langle d\theta, p \rangle + D_C(\hat{\theta}^t; \hat{\theta}^t + d\theta) + \langle \nabla C(\hat{\theta}^t + d\theta), -d\theta \rangle \\ &\geq \langle d\theta, p - \nabla C(\hat{\theta}^t + d\theta) \rangle \\ &\geq \langle d\theta, p - (1 - \lambda)\hat{p}^t \rangle \\ &= \langle d\theta, p - \hat{p}^t \rangle + \langle d\theta, \lambda\hat{p}^t \rangle. \end{aligned}$$

Now with probability $1 - \gamma$, $\|\hat{p}^t\|_\infty \leq 2$ by the accuracy guarantee, and $\lambda \leq \alpha/2$. So if $p - \hat{p}^t$ has a coordinate at least 2α , then the expected profit by buying a share in that coordinate is at least $2\alpha - \lambda(2) \geq \alpha$, which is at least the transaction fee.

The accuracy guarantee of Theorem 2 goes through unchanged.

The remainder of the proof shows the worst-case loss guarantee.

To analyze the loss of the market maker, we rephrase the mechanism into an equivalent, but easier-to-analyze scenario. There is an agent called the “noise trader” who is controlled by the market maker. When each participant t arrives, t pays the transaction fee c and purchases bundle $d\theta^t$. At this point, the market state is equal to the sum of this trade and the previous market state, *i.e.* $\hat{\theta}^{t-1} + d\theta^t$. Then, the noise trader calculates $\hat{\theta}^t$ and purchases shares from the market maker in order to move the market state to $\hat{\theta}^t$, *i.e.* the noise trader purchases a trade equal to $\hat{\theta}^t - [\hat{\theta}^{t-1} + d\theta^t]$. (The noise trader does not pay a transaction fee.) At the end of trading, the state of the world is revealed and the market maker pays out for all securities, including to the noise trader.

The loss of the market maker in the true mechanism – that is, the quantity we wish to show is bounded – is equal to the sum of the losses of the market maker and the noise trader in the

rephrased scenario. This follows because, for example, the net profit of each participating agent is identical in both scenarios.

Now, suppose that T' participants choose to trade, with $1 \leq T' \leq T$.⁷ Then, with a price-sensitivity λ , the worst-case loss is

$$WC(\lambda, T') := WC_0(\lambda, T') + NTL(\lambda, T') - T'c$$

where $WC_0(\lambda, T')$ is the worst-case loss of a standard prediction market maker with parameter λ and T' participants, $NTL(\lambda, T')$ is the worst-case noise trader loss, and $T'c$ is the revenue from T' transaction fees of size c each.

The worst-case loss of a standard prediction market maker with price sensitivity λ , by our normalization and definition of λ , is bounded by

$$WC_0(\lambda, T') \leq \frac{B}{\lambda}.$$

To bound the noise trader loss $NTL(\lambda, T')$, we will consider each “bundle” of shares z^t purchased by the noise trader. The idea is to bound the difference in price between the purchase and sale of z^t . The noise trader is then left with a small number of unsold bundles at the close of the market, which cannot have been too costly.

Observe that at a generic time step t , after participant t trades, the noise trader sells some set of previously-purchased bundles, then purchases z^t . For example, the noise traders’ purchases at the first few time steps are:

1. Sells nothing and purchases z^1 .
2. Sells z^1 and purchases z^2 .
3. Sells nothing and purchases z^3 .
4. Sells z^3 and z^2 and purchases z^4 .
5. ...

For analysis, we suppose that at each t , the noise trader first sells any previous bundles – *e.g.* at $t = 4$, first selling z^3 and then selling z^2 – and finally purchases z^t .

Now, let $b(t)$ be the largest power of 2 that divides t . Let θ_{buy}^t and θ_{sell}^t be the market state just before the noise trader purchases z^t and just after she sells z^t , respectively.

Lemma 2. *For each t , exactly $b(t)$ traders arrive between the purchase and the sale of bundle z^t ; furthermore, $\theta_{\text{sell}}^t - \theta_{\text{buy}}^t$ is exactly equal to the sum of these participants’ trades.*

For example, if t is odd, only one participant arrives between the purchase and sale of z^t ; furthermore, z^t is the last bundle purchased by the noise trader at time t and is the first sold at time $t + 1$, so the difference in market state is exactly z^t plus that participant’s trade.

⁷We wish to show that loss is bounded even if fewer than T participants arrive; otherwise, we might be in some sense relying on all T transaction fees to achieve bounded loss.

Proof of Lemma 2. Note that if we write t in binary, it has a one in the bit position $\log b(t)$, followed by zeros. By definition of the algorithm, z^t is sold at the next time $t' > t$ where the bit $\log b(t)$ is flipped to zero. So we have $t' - t = b(t)$, so $b(t)$ traders have arrived.

Now we want to show that $\theta_{\text{sell}}^t - \theta_{\text{buy}}^t$ is the sum of their trades, i.e. that every noise trade bundle held by the trader before buying z^t is still held at the moment of selling z^t , and no other noise trade bundles are held at that time.

Consider all of the noise bundles that were already held at time t (after selling the appropriate bundles at that time, but before purchasing z^t). By definition of the algorithm, these were purchased at times s with $b(s) > b(t)$, so by the above discussion, they are not sold until times t' where the bits in positions $\log b(s)$ are flipped to zero, which cannot happen until after bit $\log b(t)$ is flipped and z^t is sold. Meanwhile, every bundle purchased after z^t is sold by, at the latest, the same time that z^t is sold, as they correspond to lower-order bits; and any sold at the same time as z^t are sold first because they were purchased later. \square

Lemma 3. *If the noise trader purchases and later sells z^t , then her net loss in expectation over z^t (but for any trader behavior in response to z^t), is at most $\lambda b(t)K$ where $K = \mathbb{E} \|z^t\|_2$.*

Proof. Given the noise trader's bundle drawn is z^t , her loss is:

$$C(\theta_{\text{buy}}^t + z^t) - C(\theta_{\text{buy}}^t) + C(\theta_{\text{sell}}^t) - C(\theta_{\text{sell}}^t + z^t).$$

The first pair of terms represents the payment made to purchase z^t (moving the market state to θ_{buy}^t); the second pair represents the payment to sell z^t (moving the state to θ_{sell}^t). Lemma 2 implies that $\|\theta_{\text{buy}}^t - \theta_{\text{sell}}^t\|_1 \leq b(t)$, as each trader can buy or sell at most 1 unit of shares. Therefore, the net loss on bundle z^t is at most

$$\mathbb{E}_{z^t} \max_{\theta, \theta': \|\theta - \theta'\|_1 \leq b(t)} C(\theta + z^t) - C(\theta) + C(\theta') - C(\theta' + z^t)$$

Now, we have

$$\begin{aligned} C(\theta + z^t) - C(\theta) &= \int_{x=0}^1 \nabla C(\theta + xz^t) \cdot z^t dx, \\ C(\theta + r + z^t) - C(\theta + r) &= \int_{x=0}^1 \nabla C(\theta + r + xz^t) \cdot z^t dx. \end{aligned}$$

So the difference is

$$\begin{aligned} &\int_{x=0}^1 \langle \nabla C(\theta + xz^t) - \nabla C(\theta + r + xz^t), z^t \rangle dx \\ &\leq \int_{x=0}^1 \lambda \|r\|_2 \|z^t\|_2 dx \\ &= \lambda \|r\|_2 \|z^t\|_2 \end{aligned}$$

by definition of price sensitivity λ . We also have $\|r\|_2 \leq \|r\|_1 \leq b(t)$. This bound holds for each outcome of z^t and any behavior of the participants, so we conclude the lemma statement, that expected loss is bounded by $\lambda b(t) \mathbb{E} \|z^t\|_2$. \square

Now, for each $j = 0, \dots, \log T' - 1$, there are 2^j different steps t with $b(t) = T'/2^{j+1}$. So the total loss is at most

$$\begin{aligned} NTL(\lambda, T') &= \sum_{j=0}^{\log T' - 1} 2^j \frac{T'}{2^{j+1}} \lambda K \\ &= \frac{T' \log T'}{2} \lambda K. \end{aligned} \tag{2}$$

Now, the above assumes that the noise trader sells every bundle she buys, but in reality, she has up to $\log T'$ bundles left over at the end, which can cause an additional loss of up to $\log T' K'$ where $K' = \mathbb{E} \|z^t\|_1$, as each bundle can cost at most the sum of the shares it contains. This is insignificant compared to terms on the order of T' ; formally, we can *e.g.* bound it by (2) and therefore bound the total noise trader loss by twice (2).

This gives

$$\begin{aligned} WC(\lambda, T') &\leq WC_0(\lambda, T') + T' \log T' \lambda K - T' c \\ &\leq \frac{B}{\lambda} + T' (K \log T' \lambda - c) \\ &\leq \frac{B}{\lambda} \end{aligned} \tag{3}$$

if we choose λ and c such that $c \geq K \log T' \lambda$, in other words, $\lambda \leq c / K \log T'$.

We can bound K as follows: for each t , z^t is a vector of d independent $Lap(b)$ variables with $b = 2 \lceil \log T' \rceil / \epsilon$. By concavity of $\sqrt{\cdot}$,

$$\begin{aligned} K &= \mathbb{E} \sqrt{\sum_{i=1}^d z^t(i)^2} \\ &\leq \sqrt{\sum_i \mathbb{E} z^t(i)^2} \\ &= \sqrt{d \text{Var}(Lap(b))} \\ &= \sqrt{2db^2} \\ &= 2\sqrt{2d} \frac{\lceil \log T' \rceil}{\epsilon}. \end{aligned}$$

Therefore, it suffices to pick

$$\lambda \leq \frac{c \epsilon}{2\sqrt{2d} \lceil \log T' \rceil \log T'}.$$

For $c = \alpha$, this is already accomplished by the private, accurate market choosing λ^* as in Theorem 2. \square

Theorem 4. *For fixed $\alpha, \gamma, \epsilon, d$, there is an ϵ -differentially private prediction market using transaction fees and adaptive price-sensitivity satisfying the following:*

- *The worst case loss is a constant depending only on $\alpha, \gamma, \epsilon, d$, but not on the number of participants T .*

- The accuracy guarantee is maintained, i.e. $\|p^t - \hat{p}^t\|_1 \leq \alpha$ for all t except with probability γ .
- The incentive guarantee is maintained, i.e. a trader who believes p with $\|p - \hat{p}^t\|_\infty \geq 2\alpha$ expects to make money by participating.

Proof. Fix the parameters ϵ and d throughout. Let $\lambda^*(T, \alpha, \gamma)$ be the price sensitivity parameter as a function of these variables given in Theorem 2.

The market design is the following, with each $T^{(k)}$ to be chosen later. We run the transaction-fee private market above with $T = T^{(1)}$, transaction fee α , and price sensitivity $\lambda^{(1)} = \lambda^*(T^{(1)}, \alpha/2, \gamma/2)$. When (and if) $T^{(1)}$ participants have arrived, we create a new market whose initial state is such that its prices match the final (noisy) prices of the previous one. We set $T^{(2)}$ and price sensitivity $\lambda^{(2)} = \lambda^*(T^{(2)}, \alpha/4, \gamma/4)$ for the new market. We repeat, halving α and γ at each stage and increasing T in a manner to be specified shortly, until no more participants arrive.

ϵ -differential privacy of the market follows immediately because each stage k is differentially private for the participants who arrive in that stage, and no further information is revealed than the set of market states (the final state, which will have been revealed anyway, is used to initialize the next market).

To show the incentive guarantee, note that the transaction fee is always fixed at α , so the same incentive argument as in Theorem 3 goes through immediately.

To show the accuracy guarantee, note the the prices up to $T^{(1)}$ arrivals satisfy an $\alpha/2$ guarantee; therefore the starting prices of the new market are within $\alpha/2$ of what they would be without added noise. The prices up to $T^{(2)}$ additional arrivals are within $\alpha/2 + \alpha/4$ of what they would have been (since they begin within $\alpha/2$ and are designed to stay within $\alpha/4$ of this shifted goal); and so on, telescoping to at most α . Similarly, the chance of failure of any of these guarantees, by a union bound, is at most $\gamma/2 + \gamma/4 + \dots \leq \gamma$.

Now we must show bounded worst-case loss, and how to set $T^{(k)}$. We will choose $T^{(1)}$ to be a constant and each $T^{(k)} = 4T^{(k-1)}$.

We will claim two things:

1. In the final stage k where not all participants arrive, the market maker's loss is at most $\frac{\alpha}{16}T^{(k)}$.
2. In each stage k , if it is completed (all $T^{(k)}$ participants arrive), the market maker's profit from that stage is at least $\frac{\alpha}{2}T^{(k)}$.

These together prove bounded worst-case loss: If at least one stage is completed, the total profit is in fact positive: it is positive from all but the last stage, whose loss is at most $\frac{\alpha}{16}T^{(k)} \leq \frac{\alpha}{4}T^{(k-1)}$ which is smaller than the profit made in stage $k-1$. If no stages are completed, i.e. fewer than $T^{(1)}$ participants arrive, then loss is bounded by $\frac{1}{\lambda^{(1)}}$, which is a constant depending only on $d, \alpha, \epsilon, \gamma$.

Proof of (1): The worst-case loss in stage 1, by Theorem 3, is

$$\frac{B}{\lambda^{(1)}} = \frac{8\sqrt{2}Bd\lceil\log T^{(1)}\rceil \ln(4T^{(1)}d/\gamma)}{\alpha \epsilon} \leq \frac{\alpha}{16}T^{(1)}$$

for a large enough choice of $T^{(1)}$ as a function of $\alpha, d, \epsilon, \gamma$.⁸

⁸Fixing these parameters, the loss is upper-bounded by $C \log(T^{(1)})^2/\alpha$ for some constant C and large enough $T^{(1)}$, which is upper-bounded by $\alpha T^{(1)}/16$ for $T^{(1)}$ large enough, say $2^{C/\alpha^2}$.

Now we just show that $T^{(k)}$ increases faster than $1/\lambda^{(k)}$. $T^{(k)} = 4T^{(k-1)}$, but

$$\begin{aligned}\frac{B}{\lambda^{(k)}} &= \frac{4\sqrt{2}B2^k d \lceil \log T^{(k)} \rceil \ln(2T^{(k)}d2^k/\gamma)}{\alpha \epsilon} \\ &= 2 \frac{4\sqrt{2}B2^{k-1} d \lceil 2 + \log T^{(k-1)} \rceil (\ln(8) + \ln(2T^{(k-1)}d2^{k-1}/\gamma))}{\alpha \epsilon} \\ &\leq 4 \frac{1}{\lambda^{(k-1)}}\end{aligned}$$

for sufficiently large $T^{(k-1)}$, i.e. if $T^{(1)}$ is a sufficiently large constant. So $\frac{B}{\lambda^{(k)}}$ grows more slowly than $T^{(k)}$ and the inequality $\frac{B}{\lambda^{(k)}} \leq \frac{\alpha}{16}T^{(k)}$ continues to hold.

Proof of (2): let us lower-bound the profit in stage k if completed. By Inequality 3 (from Theorem 3), that the market-maker profit if $T' = T^{(k)}$ participants arrive is

$$\begin{aligned}&T^{(k)} \left(c - K\lambda^{(k)} \log T^{(k)} \right) - \frac{B}{\lambda^{(k)}} \\ &= T^{(k)} \left(\alpha - \frac{\sqrt{2}d \lceil \log T^{(k)} \rceil \log T^{(k)}}{\epsilon} \frac{(\alpha/2^k) \epsilon}{4\sqrt{2}d \lceil \log T^{(k)} \rceil \ln(2dT^{(k)}2^k/\gamma)} \right) - \frac{B}{\lambda^{(k)}} \\ &= \alpha T^{(k)} \left(1 - \frac{\log T^{(k)}}{4(2^k)\sqrt{d} \ln(2dT^{(k)}2^k/\gamma)} \right) - \frac{B}{\lambda^{(k)}}.\end{aligned}$$

Recall that $\frac{B}{\lambda^{(k)}} \leq \frac{\alpha}{16}T^{(k)}$. We want to conclude that the profit in stage k is at least $\frac{\alpha}{2}T^{(k)}$, so we just need to show that

$$1 - \frac{\log T^{(k)}}{4(2^k)\sqrt{d} \ln(2dT^{(k)}2^k/\gamma)} - \frac{1}{16} \geq \frac{1}{2}.$$

The fraction is decreasing in k , so it suffices to achieve this for $k = 1$, $d = 1$, and $\gamma = 1$, where we have

$$\frac{\log T^{(1)}}{8 \ln(4T^{(1)})} \leq \frac{1}{4}.$$

This suffices to prove Claim (2). □

References

- Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12, 2013. URL <http://dl.acm.org/citation.cfm?id=2465777>.
- Jacob D. Abernethy, Rafael M. Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. A General Volume-parameterized Market Making Framework. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 413–430, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2565-3. doi: 10.1145/2600057.2602900. URL <http://doi.acm.org/10.1145/2600057.2602900>.
- T-H Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)*, 14(3):26, 2011.

- Rachel Cummings, David M Pennock, and Jennifer Wortman Vaughan. The possibilities and limitations of private prediction markets. In *Proceedings of the 17th ACM Conference on Economics and Computation*, EC '16, pages 143–160. ACM, 2016.
- Cynthia Dwork and Aaron Roth. *The algorithmic foundations of differential privacy*. Foundations and Trends in Theoretical Computer Science, 2014.
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.
- Bo Waggoner, Rafael Frongillo, and Jacob D Abernethy. A Market Framework for Eliciting Private Data. In *Advances in Neural Information Processing Systems 28*, pages 3492–3500, 2015. URL <http://papers.nips.cc/paper/5995-a-market-framework-for-eliciting-private-data.pdf>.